

CLAIMS

What is claimed is:

- [c1] 1. A computer-readable medium containing a data structure for mapping between semi-structured data and structured data, the data structure comprising:
- for each parent and child relationship in the semi-structured data, a mapping from a parent identifier and a child identifier of the parent and child relationship to a table identifier of a table in the structured data that corresponds to the parent and child relationship.
- [c2] 2. The computer-readable medium of claim 1 wherein each mapping further includes a parent key identifier and a child key identifier of the table that correspond to the parent identifier and child identifier of the mapping.
- [c3] 3. The computer-readable medium of claim 2 wherein each mapping further includes an indicator of whether a child of the parent and child relationship is an attribute or an element.
- [c4] 4. The computer-readable medium of claims 1-3 wherein the data structure supports recursive definition of elements of the semi-structured data.
- [c5] 5. The computer-readable medium of claims 1-4 wherein the semi-structured data is in XML format.
- [c6] 6. The computer-readable medium of claims 1-5 wherein the structured data is in a relational format.

[c7] 7. The computer-readable medium of claims 1-6 wherein the data structure includes for each child of a parent and child relationship that is also not a parent, a mapping from the child identifier of that child to a table identifier of a table in the structured data.

[c8] 8. The computer-readable medium of claim 7 wherein the mapping from the child identifier of that child also maps to a child key identifier in the table.

[c9] 9. The computer-readable medium of claims 7-8 wherein the mapping from the child identifier of that child also maps to an indicator as to whether the child is an attribute or an element of its parent.

[c10] 10. The computer-readable medium of claims 7-9 wherein the mapping from the child identifier of that child also maps to a child value identifier in the table.

[c11] 11. A method for generating a data structure for mapping between semi-structured data and structured data, the method comprising:
identifying parent and child relationships in the semi-structured data;
for each identified parent and child relationship,
identifying a table in the structured data that
corresponds to the identified parent and child relationship; and
creating a mapping between the identified parent and child relationship and
the identified table.

[c12] 12. The method of claim 11 wherein the mapping maps a parent identifier and child identifier of the parent and child relationship to a table identifier of the identified table.

- [c13] 13. The method of claim 12 wherein each mapping further includes a parent key identifier and a child key identifier of the table that correspond to the parent identifier and child identifier of the mapping.

- [c14] 14. The method of claim 13 wherein each mapping further includes an indicator of whether a child of the parent and child relationship is an attribute or an element.

- [c15] 15. The method of claims 12-14 wherein the method supports recursive definition of elements of the semi-structured data.

- [c16] 16. The method of claims 12-15 wherein the semi-structured data is in XML format.

- [c17] 17. The method of claims 12-16 wherein the structured data is in a relational format.

- [c18] 18. The method of claims 12-17 including for each child of a parent and child relationship that is also not a parent, creating a mapping from the child identifier of that child to a table identifier of a table in the structured data.

- [c19] 19. The method of claim 18 wherein the mapping from the child identifier of that child also maps to a child key identifier in the table.

- [c20] 20. The method of claims 18-19 wherein the mapping from the child identifier of that child also maps to an indicator as to whether the child is an attribute or an element of its parent.

- [c21] 21. The method of claims 18-20 wherein the mapping from the child identifier of that child also maps to a child value identifier in the table.

[c22]

22. A method comprising:
generating a structured organization to store a collection of semi-structured data; and
collaterally generating a description of how the semi-structured data is stored under the structured organization.

[c23]

23. The method of claim 22, wherein said collateral generation of a description comprises generating descriptive entries for a description table, with each descriptive entry identifying a first and a second storage column of a storage table for a parent and a child node of a parent and child node pair of the semi-structured data.

[c24]

24. The method of claim 23, wherein said collateral generation of a description further comprises generating for each of a selected one or ones of said descriptive entries a context qualifier for at least a selected one of the parent and the child node.

[c25]

25. The method of claim 22, wherein said generation of a structured organization comprises creating relational storage tables with rows and columns in a normalized manner.

[c26]

26. An apparatus comprising:
a storage medium having stored therein a plurality of programming instructions designed to generate a structured organization to store a collection of semi-structured data, when executed, and to collaterally generate a description of how the semi-structured data are stored under the structured organization; and a processor coupled to the storage medium to execute the programming instructions.

[c27] 27. The apparatus of claim 26, wherein the programming instructions are designed to generate the description by generating a description table having a plurality of descriptive entries, with each descriptive entry identifying a first and a second storage column of a storage table of the structured organization for a parent and a child node of a parent and child node pair of the semi-structured data.

[c28] 28. The apparatus of claim 27, wherein said programming instructions are further designed to generate for each of a selected one or ones of said descriptive entries a context qualifier for at least a selected one of the parent and the child node.

[c29] 29. The apparatus of claim 26, wherein said programming instructions are designed to generate a structured organization by creating relational storage tables with rows and columns in a normalized manner.

[c30] 30. The apparatus of claim 26, wherein said programming instructions are a subset of a larger collection of programming instructions implementing a selected one of a database manager and a development toolkit that facilitates development of applications that access databases.

[c31] 31. An article of manufacture comprising:
a recordable medium having recorded thereon a plurality of programming instructions for use to program a system having a processor for executing programming instructions, to enable the system to generate a structured organization to store a collection of semi-structured data, and to collaterally generate a description of how the semi-structured data are stored under the structured organization.

[c32] 32. The article of claim 31, wherein said programming instructions enable the system to collaterally generating a description by generating a description table having a plurality of descriptive entries, with each descriptive entry identifying a first and a second storage column of a storage table for a parent and a child node of a parent and child node pair of the semi-structured data.

[c33] 33. The article of claim 32, wherein said programming instructions further enable the system to generate for each of a selected one or ones of said descriptive entries a context qualifier for at least a selected one of the parent and the child node.

[c34] 34. The article of claim 31, wherein said programming instructions enable the system to generate a structured organization by creating relational storage tables with rows and columns in a normalized manner.

[c35] 35. The article of claim 31, wherein said programming instructions are a subset of a larger collection of programming instructions implementing a selected one of a database manager and a development toolkit that facilitates development of applications that access databases.

[c36] 36. A method comprising:
generating a semi-structured data organization for a collection of structured data; and
collaterally generating a description of correspondence between the generated semi-structured data organization and the collection of structured data.

[c37] 37. The method of claim 36, wherein said collateral generation of a description comprises generating descriptive entries for a description table, with

each descriptive entry identifying a first and a second storage column of a storage table of the structured data for a parent and a child node of a parent and child node pair of the generated semi-structured data organization.

[c38] 38. The method of claim 37, wherein said collateral generation of a description further comprises generating for each of a selected one or ones of said descriptive entries a context qualifier for at least a selected one of the parent and the child node.

[c39] 39. The method of claim 37, wherein said generation of a semi-structured data organization comprises creating an extensible mark-up language (XML) data structure for a collection of relational tables.

[c40] 40. An apparatus comprising: a storage medium having stored therein a plurality of programming instructions designed to generate a semi-structured data organization for a collection of structured data, when executed, and to collaterally generate a description of correspondence between the semi-structured data organization and the structured data; and a processor coupled to the storage medium to execute the programming instructions.

[c41] 41. The apparatus of claim 40, wherein the programming instructions are designed to generate the description by generating a description table having a plurality of descriptive entries, with each descriptive entry identifying a first and a second storage column of a storage table of the structured data for a parent and a child node of a parent and child node pair of the generated semi-structured data organization.

[c42] 42. The apparatus of claim 41, wherein said programming instructions are further designed to generate for each of a selected one or ones of said

descriptive entries a context qualifier for at least a selected one of the parent and the child node.

[c43] 43. The apparatus of claim 40, wherein said programming instructions are designed to create an extensible mark-up language (XML) data structure for a collection of relational tables.

[c44] 44. The apparatus of claim 40, wherein said programming instructions are a subset of a larger collection of programming instructions implementing a selected one of a database manager and a development toolkit that facilitates development of applications that access databases.

[c45] 45. An article of manufacture comprising:
a recordable medium having recorded thereon a plurality of programming instructions for use to program a system having a processor for executing programming instructions, to enable the system to generate a semi-structured data organization for a collection of structured data, and to collaterally generate a description of correspondence between the semi-structured data organization and the structured data.

[c46] 46. The article of claim 45, wherein said programming instructions enable the system to collaterally generating a description by generating a description table having a plurality of descriptive entries, with each descriptive entry identifying a first and a second storage column of a storage table of the structured data for a parent and a child node of a parent and child node pair of the generated semi-structured data organization.

[c47] 47. The article of claim 46, wherein said programming instructions further enable the system to generate for each of a selected one or ones of said

descriptive entries a context qualifier for at least a selected one of the parent and the child node.

[c48] 48. The article of claim 45, wherein said programming instructions enable the system to create an extensible mark-up language (XML) data structure for a collection of relational tables.

[c49] 49. The article of claim 45, wherein said programming instructions are a subset of a larger collection of programming instructions implementing a selected one of a database manager and a development toolkit that facilitates development of applications that access databases.

[c50] 50. A method comprising:
pre-processing semi-structured queries in an information page employing semi-structured data to identify the data to be selectively retrieved from the semi-structured data, and the storage locations of the data in a collection of structured data comprising the semi-structured data, referencing a description of how the semi-structured data are stored in the collection of structured data;
generating equivalent structured queries to retrieve the data from the collection of structured data; and
replacing the semi-structured queries with the equivalent structured queries.

[c51] 51. The method of claim 50, wherein said pre-processing of semi-structured queries is also performed to identify join conditions for joining relational storage tables of said collection of structured data to retrieve the data.

[c52] 52. The method of claim 50, wherein said pre-processing of semi-structured queries comprises pre-processing matching templates of said semi-structured queries.

[c53] 53. The method of claim 52, wherein said pre-processing matching templates comprises pre-processing of data extraction commands within each of the matching templates.

[c54] 54. The method of claim 50, wherein said generation of equivalent structured queries comprises generating structured queries, and control structures controlling repeated execution of the generated structured queries.

[c55] 55. An apparatus comprising:
a storage medium having stored therein a plurality of programming instructions designed
to pre-process semi-structured queries in an information page employing semi-structured data to identify the data to be selectively retrieved from the semi-structured data, and the storage locations of the data in a collection of structured data comprising the semi-structured data, referencing a description of how the semi-structured data are stored in the collection of structured data,
to generate equivalent structured queries to retrieve the data from the collection of structured data, and
to replace the semi-structured queries with the equivalent structured queries; and
a processor coupled to the storage medium to execute the programming instructions.

[c56] 56. The apparatus of claim 55, wherein said programming instructions are also designed to pre-process the semi-structured queries to identify join conditions for joining relational storage tables of said collection of structured data to retrieve the data.

[c57] 57. The apparatus of claim 55, wherein said programming instructions are designed to pre-process matching templates of said semi-structured queries.

[c58] 58. The apparatus of claim 57, wherein said programming instructions are further designed to pre-process data extraction commands within each of the matching templates.

[c59] 59. The apparatus of claim 55, wherein said programming instructions are designed to generate structured queries, and control structures controlling repeated execution of the generated structured queries.

[c60] 60. The apparatus of claim 55, wherein said programming instructions are a subset of a larger collection of programming instructions implementing a selected one of a database manager and a development toolkit that facilitates development of applications that access data bases.

[c61] 61. An article of manufacture comprising:
a recordable medium having recorded thereon a plurality of programming instructions for use to program a system having a processor for executing programming instructions, to enable the system
to pre-process semi-structured queries in an
information page employing semi-structured data to identify the data to be selectively retrieved from the semi-structured data, and the storage locations of the data in a collection of structured data comprising the semi-structured data, referencing a description of

how the semi-structured data are stored in the collection of structured data,

to generate equivalent structured queries to retrieve the data from the collection of structured data, and

to replace the semi-structured queries with the equivalent structured queries.

[c62] 62. The article of claim 61, wherein said programming instructions are also designed to enable the system to pre-process the semi-structured queries to identify join conditions for joining relational storage tables of said collection of structured data to retrieve the data.

[c63] 63. The article of claim 61, wherein said programming instructions are designed to enable the system to pre-process matching templates of said semi-structured queries.

[c64] 64. The article of claim 63, wherein said programming instructions are further designed to enable the system to pre-process data extraction commands within each of the matching templates.

[c65] 65. The article of claim 61, wherein said programming instructions are designed to enable the system to generate structured queries, and control structures controlling repeated execution of the generated structured queries.

[c66] 66. The article of claim 61, wherein said programming instructions are a subset of a larger collection of programming instructions implementing a selected one of a database manager and a development toolkit that facilitates development of applications that access data bases.

[c67]

67. A method comprising:

storing pre-processed information pages formed with a language that provides for semi-structured queries, and employing semi-structured data, where the semi-structured queries have been pre-replaced with equivalent structured queries to retrieve the data to be retrieved by the semi-structured queries from the semi-structured data, from a collection of structured data comprising the semi-structured data instead;

receiving a request for a selected one of the information pages;

dynamically completing the information page, executing the replacement structured queries; and

returning the dynamically completed information page.

[c68]

68. The method of claim 67, wherein the method further comprises pre-processing the semi-structured queries of each information page

employing semi-structured data to identify the data to be selectively retrieved from the semi-structured data, and the storage locations of the data in the collection of structured data comprising the semi-structured data, referencing a description of how the semi-structured data are stored in the collection of structured data;

generating equivalent structured queries to retrieve the data from the collection of structured data; and

replacing the semi-structured queries with the equivalent structured queries.

[c69]

69. The method of claim 68, wherein said pre-processing of semi-structured queries is also performed to identify join conditions for joining relational storage tables of the collection of structured data to retrieve the data.

[c70] 70. The method of claim 68, wherein said generation of equivalent structured queries comprises generating structured queries, and control structures controlling repeated execution of the generated structured queries.

[c71] 71. An apparatus comprising:
a network interface coupling the apparatus to a data network;
a storage medium having stored therein pre-processed information pages formed with a language that provides for semi-structured queries, and employing semi-structured data, where the semi-structured queries have been pre-replaced with equivalent structured queries to retrieve the data to be retrieved by the semi-structured queries from the semi-structured data, from a collection of structured data comprising the semi-structured data instead; and
a processor coupled to the network interface and the storage medium to receive a request from a client for a selected one of the information pages from the data network, and in response, dynamically completing the information page, executing the replacement structured queries, and returning the dynamically completed information page to the client through the data network.

[c72] 72. The apparatus of claim 71, wherein the apparatus is a web server, the language is extensible script language (XSL), the information pages are XSL documents, and the semi-structured data are extensible mark-up language (XML) data.

[c73] 73. The apparatus of claim 72, wherein the structured data are relational data and the structured queries are formed using the structured query language (SQL).